

## HIGH PERFORMANCE PACKET PROCESSING USING A GENERAL PURPOSE PROCESSOR

5

### BACKGROUND OF THE INVENTION

The field of the invention is computer network security, and more specifically a system for rapidly identifying and processing incoming packets from a computer network.

Many computer network security tools require a high performance packet-processing capability to identify, classify and act upon on incoming data packets according to predefined rules. Over the years, increasing network carrying capacity and speed, and increasing use of the Internet, have resulted in a large increase in the number of data packets routed into systems connected to data networks. Further, the average degree of complexity of actions that are performed on those data packets has increased, due to the increased number of security threats that may be carried over a computer network. For example, distributed denial of service (DDOS) attacks flood one or more computers with a vast number of packets, such that the target computer cannot effectively process other data, blocking access to those one or more computers by individuals or entities wishing to connect legitimately to that target computer. The recent spread of such DDOS attacks on well-known companies through the World Wide Web shut down many e-commerce web sites and network business activities, and highlighted the potential security threats posed by network connections.

Existing packet-processing systems typically deal with packet complexity and volume in an ad hoc way, and rely on one or more complicated application specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs) dedicated to packet processing, thereby providing a fast but inflexible solution. A general processor may be integrated with the dedicated packet processing ASICs or FPGAs in an attempt to provide flexibility to the system. However, such integration generally provides poor results, degrading the overall performance of the packet processing system. Further, existing packet processing systems typically use a time-

expensive hashing mechanism to assign rules to packets and process them, further degrading packet processing performance.

Thus, there is a need for a system and method for providing efficient integration of processors for data packet processing. Further, there is a need for a system and method for 5 providing more efficient complex processing of data packets routed into computer systems. Further, there is a need for a system and method for organizing complex processing of data packets.

#### SUMMARY OF THE INVENTION

The present invention satisfies the shortcomings of data packet processing systems in the prior art. The system of the present invention comprises a general-purpose high-performance packet processor. A policy action table is implemented to support the flexibility of packet processing operations and fast network policy changes.

In an aspect of a preferred embodiment, a packet processing device includes a control logic processor for filtering packets according to a set of stored rules and an arithmetic logic processor for executing packet processing instructions based on the content of the packet. In a further aspect of a preferred embodiment, the control logic processor spawns a new thread for each incoming packet, relieving the arithmetic logic processor of the need to do so. In a further aspect of a preferred embodiment, the control logic processor assigns a policy to each incoming 20 packet.

In another aspect of a preferred embodiment, a policy action table stores one or more policy instructions. In a further aspect of a preferred embodiment, the policy action table provides a high degree of flexibility by allowing policy changes through an update of the entries in the policy action table. In a further aspect of a preferred embodiment, the policy action table 25 maps a virtual packet flow identification code to the physical memory address of an action code and a state block associated to the identification code. In a further aspect of a preferred

embodiment, the policy action table is implemented with a general-purpose processor to deliver a high degree of flexibility and performance.

In another aspect of a preferred embodiment, the arithmetic logic processor processes a packet based on the stored policy assigned to that packet.

5

#### BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating a system architecture for present invention.

Fig. 2 is a flow diagram illustrating a method performed by the system of Fig. 1.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring to FIG. 1, a preferred embodiment of a packet processor 2 is shown. There are two main functional units: an arithmetic logic unit (ALU) 20 and a control logic unit (CLU) 10. The ALU 20 is the processor core and the CLU 10 is the filtering unit. In a preferred embodiment, the CLU 10 analyzes packets accepted by a register-in 4 and determines whether to accept or discard those packets. Such analysis is preferably performed by logic 12 within the CLU 10, as shown schematically in FIG. 1. In a preferred embodiment, the CLU 10 is implemented on an application-specific integrated circuit (ASIC) or field-programmable gate array (FPGA), both of which are well known to those skilled in the art. The CLU 10 assigns a thread to each accepted packet, which is appended to or otherwise associated with those accepted packets. Preferably, the ALU 20 processes the accepted packets based on a policy to which those packets are assigned, and transmits such processed packets to a register-out 3. Such processed packets then proceed to their destination. The ALU 20 is preferably a general-purpose microprocessor, and may be a complex instruction set computer processor, or a reduced instruction set computer processor. The packet processing queue 150, which holds threads associated with packets admitted by the CLU 10 for later processing by the ALU 20, is preferably implemented in a computer memory such as a random-access memory (RAM) chip,

which is commonly available and well known to those skilled in the art. However, other forms of memory may be used, such as but not limited to flash memory. The policy action table 176 maps the threads associated with the packets that are stored in the packet processing queue 150 to the physical memory address of the instructions and state associated with each packet. In a 5 preferred embodiment, the policy action table 176 is implemented in an ASIC, allowing fast access between the policy action table 176 and the thread data it acts upon, and allowing the instructions in the policy action table 176 to be easily revised. The policy action table 176 may be implemented in structures other than an ASIC, if desired. The memory unit 100 preferably takes the form of random-access memory (RAM) chips, which are commonly available and well 10 known to those skilled in the art. However, other forms of memory may be used, such as but not limited to flash memory. The memory unit 100 preferably stores packet processing instructions and state data. The memory unit 100 may be implemented in the same physical memory device as the packet processing queue 150, or in a physically separate memory device. The packet processing queue 150 may also be implemented in a processor cache associated with the ALU 20. The use of a processor cache is well known to those skilled in the art. A data bus 30 15 connects the memory unit to the CLU 10 and the ALU 20. Preferably, the data bus 30 operates at a speed fast enough to allow packet processing to occur rapidly enough such that incoming packets do not overflow the packet processing queue 150. Several registers and counters are also shown in FIG. 1, the functions of which are described in greater detail below. The physical 20 implementation of registers and counters is well known to those skilled in the art.

Referring as well to FIG. 2, a preferred embodiment of a method for packet processing is shown. In step 200, an incoming packet is received into the register-in 4. The incoming packet is preferably received from a communications network such as the Internet, but may instead be received from a local area network or other type of network, or from another machine in another 25 fashion, such as but not limited to wireless transmission. The assembly and transmission of data packets is well known to those of ordinary skill in the art. The incoming packet is held

temporarily in the register-in 4 until it is processed by the CLU 10. It is expected that a number of packets will be received by the register-in 4 and held there for processing. In a preferred embodiment, the register-in 4 is sized to accommodate the maximum expected packet traffic, to minimize the risk that packets transmitted to the packet processor 2 will be inadvertently lost.

5 Next, in step 202, a packet stored in the register-in 4 is transmitted to the CLU 10. Preferably, the register-in 4 transmits packets to the CLU 10 on a first-in, first-out (FIFO) basis, because the characteristics of the packets stored in the register-in 4 that would be used to determine priority have yet to be determined. A policy condition table 50 is associated with the CLU 10. The policy condition table 50 contains one or more predefined rules for admitting incoming packets from the register-in 4. A first rule 52 and a second rule 54 are shown schematically within the policy condition table 50 as examples; in a preferred embodiment, many more rules may be stored within the policy condition table 50. The rules within the policy condition table may correspond to packet formats or packet contents. The CLU 10 reads the packet header information and/or the packet contents, and searches for a matching rule in the policy condition table 50.

10 Next, in step 204, three possible outcomes of that search are possible. First, if the incoming packet does not match any rule stored in the policy condition table 50, the process moves to step 206, in which the packet is discarded. By discarding nonstandard packets of unexpected types, the CLU 12 improves security. Such nonstandard packets are likely to be hostile in intent and effect, and preventing them from entering a computer or network and being processed enhances network security. Second, if the incoming packet matches a rule instructing the CLU 10 to discard the incoming packet, the process proceeds to step 206, where the incoming packet is discarded. A rule instructing the CLU 10 to discard a particular type of packet is particularly useful when particular repeating packets are known to be hostile. For example, in a typical DDOS attack a vast number of identical packets are transmitted to the target computer to consume most or all of the target computer's processing time and/or other

system resources. A rule instructing the CLU 10 to discard such packets allows a computer or network connected to the packet processor 2 to function substantially normally, because those hostile packets are not allowed into the computer or network to consume valuable processing time. Finally, if the incoming packet matches a rule instructing the CLU 10 to forward the 5 packet for processing, the process advances to step 208.

In step 208, the CLU 10 places the filtered incoming packet 122 into the packet processing queue 150, and generates a thread 232 corresponding to the packet 122. A thread is a logic identifier 122 or other pointer associated with a record 178 stored in the policy action table 176. The logic identifier 232 is appended to or otherwise associated with the packet 122 for which it was generated. For example, if the incoming packet matches the first rule 52 in the policy condition table 50, the CLU 10 generates a thread 232 which points to a record 178 stored in the policy action table 176 for processing a packet in accordance with the first rule 52, and appends that thread to the packet 122. That is, the thread 232 is matched to a record 178 in the policy action table 176 that corresponds to the rule to which the incoming packet 122 was matched by the CLU 10.

The policy action table 176 contains one or more policies for packet processing, corresponding to the rules contained within the policy condition table 50. The policy condition table 50 assigns a rule to a packet; the policy action table 176 acts on the corresponding policy 178 and directs the processing of the packet according to that policy. Network policy is not limited to concerns such as security, and may include policies such as resource usage priority and bandwidth management. The policy action table 176 provides a high degree of policy flexibility. For example, to change a policy regarding a certain type of packet, preferably a single change to the policy action table 176 will suffice. Packets corresponding to that new policy are then processed according to its instructions.

Also in step 208, state data derived from the packet header or packet contents preferably is loaded into the fourth register 146, which buffers such state data before writing it to the

memory unit 100. Preferably, this state data is used to generate one or more state blocks stored in the memory unit 100. The use of one or more state blocks in packet processing is described in greater detail below.

In step 210, the packet 122 and its associated thread 232 preferably are placed in the 5 packet processing queue 150, which is a temporary storage location for packets to be processed. Next, in step 212, that packet and its associated thread wait in the packet processing queue 150 for processing by the ALU 20. The packet processing queue 150 may be ordered as a FIFO queue, a priority queue, or other form of queue. In a FIFO queue, packets entering the packet processing queue 150 are processed in the order in which they enter the packet processing queue 10 150. In a priority queue, packets entering the packet processing queue 150 are processed in an order based on one or more of their characteristics. For example, packets addressed to a certain machine on a network, or having certain content, may be processed before other packets entering the packet processing queue 150. Preferably, steps 210 and 212 are omitted from the process if no other packets are in the packet processing queue 150 and the ALU 20 is available for packet processing. However, step 210 may still be utilized, by placing packets in the packet processing queue 150 and taking them up immediately for processing. No artificial waiting period need be imposed on a packet in the packet processing queue 150.

In step 214, when the ALU 20 is available for processing the next packet awaiting processing in the packet processing queue 150, the program counter register 180, the first register 20 140, the second register 142, and the third register 144 are loaded. The logic identifier 232 of that packet 122 is loaded into the program counter register 180. The value of the logic identifier 232 is in turn matched with the corresponding record 178 in the policy action table 176, such that the address pointer 130 associated with a particular packet is read from the record 178 to indicate 25 that the first instruction for processing the thread is stored at the location 130 in the memory unit 100. Preferably, if additional locations in the memory unit 100 are required for storing processing instructions, the first location 130 includes an address for proceeding to another

location in the memory unit 100 for further instructions. A particular policy 178 to be executed on a packet contains an instruction address pointer 130 that points to a location 102 in the memory unit 100 where the instructions corresponding to the policy 178 are stored, and a state address pointer 132 that points to a location 104 in the memory unit 100 where a state block is stored. In a preferred embodiment, the state block includes one or more data items corresponding to the context (that is, the state of the machine or the state of the network) with which the packet processor 2 is used. Preferably, the instructions 102 for processing a packet take into account one or more elements of the state of the machine or the state of the network that will utilize that packet. In a preferred embodiment, multiple state blocks are present within the memory unit 100, each containing specific elements of the machine or network state for use by the associated instruction 102 in packed processing. However, a single state block may be provided for use by each of the instructions 102 in the memory unit 100. The state block or blocks are preferably dynamically updated to reflect the current state of the machine or network. By adjusting packet processing based on machine state, some network policies may be better implemented, such as but not limited to policies regarding bandwidth management and congestion control. The state block is preferably updated based on state data received from incoming packet headers via the fourth register 146. The determination of a network state and updating the state of a network connection are well known to those in the art.

The first register 140, the second register 142 and the third register 144 are loaded according to the policy action table 176. When the policy 178 is to be executed, the instructions 102 for executing that policy are loaded from the memory unit 100 into the first register 140, and the packet itself is loaded from the packet processing queue 150 into the second register 142. The address pointer element 132 is loaded into the third register 144 to specify a base address 104 of the associated state block in the memory unit 100, such that the instructions 102 may utilize the state block into account in processing the thread if the instructions 102 so require. However, the registers may be used in other ways if desired.

In step 216, the instructions 102 are executed. The ALU 20 reads the instructions 102 stored in the first register 140, and uses them to operate upon the packet stored in the second register 142. The instructions 102 located at the instruction pointer location 130 for the thread 232 preferably are loaded and executed by the ALU 20 in sequence. If the ALU 20 comprises 5 more than one processor, the instructions may be loaded and executed in series or in parallel.

The ALU 20, as a part of its core processing duties, provides a slice of time for processing threads waiting in the packet processing queue 150. When processing a thread, the ALU 20 locates a record 178 in the policy action table 176 corresponding to the logic identifier 232. The logic identifier 232 has a value according to the rule that the packet matched with 10 record 178. The record 178 contains a first memory address pointer 130 indicating the beginning of a code segment or set of executable instructions 102 for executing on the thread, and a second memory address pointer 132 indicating the location of the state block associated with that set of instructions 102. The ALU 20 thus may process multiple threads according to respective sets of instructions located in the memory unit 100.

After each instruction from instruction set 102 is loaded, the program counter 180 is updated by the program counter increment unit 182 as long as a get next thread or end of 15 instruction set indicator is not encountered by the ALU 20. The get next thread and end of instruction set indicators are preferably placed in the instruction set 102 to inform the ALU 20 when the instruction set 102 for a packet has been completely executed, so that the ALU 20 can 20 proceed to the next packet and expeditiously continue processing packets waiting in the packet processing queue 150.

Finally, in step 218 the processed packet is transmitted to the register-out 3, or a table or other data storage entity, for use or further processing. The register-out 3 is preferably a standard 25 register or buffer, capable of storing enough processed packets that the risk of loss of a processed packet is minimized.

The packet processor 2 is capable of processing a number of different incoming packets.

These incoming packets are either discarded or admitted, as described above. Each admitted packet is assigned a thread by the CLU 10 and placed in the packet processing queue 150, also as disclosed above. As an example, when a second packet 124 is admitted by the CLU 10, it is assigned a thread 234. The second thread 234 corresponds to a specific policy within the policy action table 176. The second thread 234 need not point to a different policy than the first thread 232; indeed, if the first packet 122 and the second packet 124 both have been classified according to the same rule in step 202, both the first thread 232 and the second thread 234 will point to the same policy 178. However, if the first packet 122 and the second packet 124 are classified differently in step 202, then the first thread 232 preferably points to a different policy in the policy action table 176 than the second thread 234. If the packet processing queue 150 is a FIFO queue, the first packet 122 is processed by the ALU 20 before the second packet 124, which awaits its turn in the packet processing queue 150. If the packet processing queue 150 is a priority queue, the second packet 124 may be processed ahead of the first packet 122 based on their relative priority. The processing of the second packet 124, as well as subsequent packets admitted by the packet processor 2, proceeds as disclosed above with regard to the first packet 122.

More than one CLU 10 and/or ALU 20 may be utilized if desired. The use of multiple processors may enhance the speed with which the packet processor 2 filters and processes incoming packets. If more than one CLU 10 is used, the size of the packet processing queue 150 may be adjusted to handle the increased rate of packet intake. Further, if multiple CLUs 10 or ALUs 20 are used, then the size of the memory unit 100 also may be adjusted accordingly, as may the capacity and the number of registers.

While the preferred embodiment has been described in terms of the creation and utilization of a number of threads corresponding to different packets, it is within the scope of the preferred embodiment to associate different packets to different tokens or flows instead of threads, and process those tokens or flows instead of threads. The use and processing of tokens

and data flows are well known to those skilled in the art.

A preferred apparatus and method for high performance packet processing and many of their attendant advantages has thus been disclosed. It will be apparent, however, that various changes may be made in the form, construction and arrangement of the parts or in the steps or  
5 order of steps of the method, without departing from the spirit and scope of the invention, the form hereinbefore described being merely a preferred or exemplary embodiment thereof. Therefore, the invention is not to be restricted or limited except in accordance with the following claims and their legal equivalents.

00000000-0000-0000-0000-000000000000